

**IN THE CLAIMS:**

- 1 1. (ORIGINAL) A method for uniformly distributing data transmitted by a server over a  
2 plurality of underlying links of an aggregate within a computer network, the method com-  
3 prising the steps of:
  - 4       defining a unit of data as a datagram;
  - 5       apportioning each datagram into at least one fragment at the server;
  - 6       associating each fragment to an underlying link of the aggregate on the basis of an  
7 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the  
8 aggregate; and
  - 9       transmitting the fragment over its associated underlying link from the server to the  
10 computer network.
- 1 2. (ORIGINAL) The method of Claim 1 wherein the step of associating comprises the  
2 step of producing a result representing a remainder upon dividing the IP ID by the num-  
3 ber of active links.
- 1 3. (ORIGINAL) The method of Claim 2 wherein the step of associating further comprises  
2 the steps of:
  - 3       calculating the IP ID of each datagram in a sequential manner; and
  - 4       rotating the fragments of each datagram among all the underlying links to thereby  
5 ensure that all fragments having the same IP ID are provided to the same physical link of  
6 the aggregate.

1 4. (ORIGINAL) The method of Claim 1 wherein the step of associating comprises the  
2 steps of:

3 logically combining the IP ID with a predetermined mask to produce a quantity;

4 right shifting the quantity a predetermined number of places; and

5 establishing a threshold at which a group of data is forwarded to each underlying  
6 link of the aggregate.

1 5. (ORIGINAL) The method of Claim 4 wherein the step of associating further comprises  
2 the step of producing a result representing a remainder upon dividing the right shifted  
3 logically combined quantity IP ID and predetermined mask by the number of active links.

1 6. (ORIGINAL) The method of Claim 5 wherein the IP ID is a 16-bit value, the prede-  
2 termined mask is 0xFF80 and predetermined number of right shifted places is 7, and  
3 wherein the group of data comprises 128 IP IDs.

1 7. (ORIGINAL) The method of Claim 6 wherein the group of data comprises one of 128  
2 different transport control protocol (TCP) fragments and 128 different user datagram pro-  
3 tocol (UDP) datagrams.

1 8. (ORIGINAL) The method of Claim 7 wherein each UDP datagram comprises up to 23  
2 fragments.

1 9. (ORIGINAL) The method of Claim 1 further comprising the steps of:

2 loading at least one data buffer of the server with the at least one fragment;

3 fetching the fragment from the data buffer; and

4 loading at least one queue of the server with the fragment, the queue associated  
5 with the underlying link.

1 10. (ORIGINAL) A system adapted to uniformly distributing data over a plurality of un-  
2 derlying links of an aggregate within a computer network, the system comprising:

3 a processor;

4 a memory coupled to the processor and having locations addressable by the proc-  
5 essor;

6 an operating system resident in the memory locations and executed by the proces-  
7 sor, the operating system configured to implement a modified load balancing technique  
8 that defines a unit of data as a datagram, the operating system comprising an Internet Pro-  
9 tocol (IP) layer that apportions the datagram into at least one fragment, the operating sys-  
10 tem further comprising a virtual interface process that associates the fragment to an un-  
11 derlying link of the aggregate on the basis of an IP identifier (ID) of the datagram and a  
12 number of active links of the aggregate; and

13 at least one network adapter coupled to the memory and processor that cooperates  
14 with a network driver of the operating system to transmit the fragment over the associated  
15 underlying link to the computer network.

1 11. (ORIGINAL) Apparatus for uniformly distributing data transmitted by a server over a  
2 plurality of underlying links of an aggregate within a computer network, the apparatus  
3 comprising:

4 means for defining a unit of data as a datagram;

5 means for apportioning each datagram into at least one fragment at the server;

6 means for associating each fragment to an underlying link of the aggregate on the  
7 basis of an Internet protocol (IP) identifier (ID) of each datagram and a number of active  
8 links of the aggregate; and

9 means for transmitting the fragment over its associated underlying link from the  
10 server to the computer network.

1 12. (ORIGINAL) The apparatus of Claim 11 wherein the means for associating com-  
2 prises means for producing a result representing a remainder upon dividing the IP ID by  
3 the number of active links.

1 13. (ORIGINAL) The apparatus of Claim 12 wherein the means for associating further  
2 comprises:

3 means for calculating the IP ID of each datagram in a sequential manner; and

4 means for rotating the fragments of each datagram among all the underlying links  
5 to thereby ensure that all fragments having the same IP ID are provided to the same  
6 physical link of the aggregate.

1 14. (ORIGINAL) The apparatus of Claim 11 wherein the means for associating com-  
2 prises:

3 means for logically combining the IP ID with a predetermined mask to produce a  
4 quantity;

5 means for right shifting the quantity a predetermined number of places; and

6 means for establishing a threshold at which a group of data is forwarded to each  
7 underlying link of the aggregate.

1 15. (ORIGINAL) The apparatus of Claim 14 wherein the means for associating further  
2 comprises means for producing a result representing a remainder upon dividing the right  
3 shifted logically combined quantity IP ID and predetermined mask by the number of ac-  
4 tive links.

1 16. (ORIGINAL) A computer readable medium containing executable program instruc-  
2 tions for uniformly distributing data transmitted by a server over a plurality of underlying  
3 links of an aggregate within a computer network, the executable program instructions  
4 comprising program instructions for:

5 defining a unit of data as a datagram;

6 apportioning each datagram into at least one fragment at the server;

7 associating each fragment to an underlying link of the aggregate on the basis of an  
8 Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the  
9 aggregate; and

10 transmitting the fragment over its associated underlying link from the server to the  
11 computer network.

1 17. (ORIGINAL) The computer readable medium of Claim 16 wherein the program in-  
2 struction for associating comprises a program instruction for producing a result represent-  
3 ing a remainder upon dividing the IP ID by the number of active links.

1 18. (ORIGINAL) The computer readable medium of Claim 17 wherein the program in-  
2 struction for associating further comprises program instructions for:

3 calculating the IP ID of each datagram in a sequential manner; and

4 rotating the fragments of each datagram among all the underlying links to thereby  
5 ensure that all fragments having the same IP ID are provided to the same physical link of  
6 the aggregate.

1 19. (ORIGINAL) The computer readable medium of Claim 16 wherein the program in-  
2 struction for associating comprises program instructions for:

3 logically combining the IP ID with a predetermined mask to produce a quantity;

4 right shifting the quantity a predetermined number of places; and

5 establishing a threshold at which a group of data is forwarded to each underlying link of  
6 the aggregate.

1 20. (ORIGINAL) The computer readable medium of Claim 19 wherein the program in-  
2 struction for associating further comprises the program instruction for producing a result  
3 representing a remainder upon dividing the right shifted logically combined quantity IP  
4 ID and predetermined mask by the number of active links

1 21. (PREVIOUSLY PRESENTED) A method for distributing data over a plurality of  
2 network links within a computer network, comprising the steps of:

3 defining a unit of data as a datagram;

4 apportioning each datagram into at least one fragment;

5 associating each fragment to a network link of the plurality of network links ac-  
6 cording to a round robin policy based at least in part on an Internet protocol (IP) identifier  
7 (ID) of each datagram;

8 transmitting the fragment over the fragment's associated network link.

1 22. (PREVIOUSLY PRESENTED) The method of claim 21 wherein the step of associ-  
2 ating is further based, at least in part, on a number of network links in the plurality of  
3 links.

1 23. (PREVIOUSLY PRESENTED) The method of claim 22 wherein the step of associ-  
2 ating comprises the step of:

3 producing a result representing a remainder by dividing the IP ID by the number  
4 of network links.

1 24. (PREVIOUSLY PRESENTED) The method of claim 23 wherein the step of associ-  
2 ating further comprises the steps of:

3 calculating the IP ID of each datagram in a sequential manner; and

4 rotating the fragments of each datagram among all the network links of the plural-  
5 ity of network links to thereby ensure that all fragments having the same IP ID are pro-  
6 vided to the same network link.

1 25. (PREVIOUSLY PRESENTED) The method of claim 21 wherein the step of associ-  
2 ating comprises the steps of:

3 logically combining the IP ID with a predetermined mask to produce a quantity;

4 right shifting the quantity by a predetermined number of places to create a new  
5 quantity; and

6 establishing a threshold at which a group of data is forwarded to each network  
7 link of the plurality of network links.

1 26. (PREVIOUSLY PRESENTED) The method of claim 25 wherein the step of associ-  
2 ating further comprises the step of producing a result representing a remainder upon di-  
3 viding the new quantity by a number of network links in the plurality of network links.

1 27. (CURRENTLY AMENDED) A system for distributing data over a plurality of net-  
2 work links within a computer network comprising:

3 a processor;

4 a memory coupled to the processor and having locations addressable by the proc-  
5 essor;

6 an operating system resident in the memory locations and executed by the proces-  
7 sor, the operating system configured to implement a load balancing technique that defines  
8 a unit of data as a datagram, the operating system further configured to include ~~including~~  
9 an Internet Protocol (IP) layer that apportions the datagram into at least one fragment, the  
10 operating system further configured to include ~~including~~ a virtual interface process that  
11 associates each fragment to a network link of the plurality of network links according to a  
12 round robin policy based at least in part on an Internet protocol (IP) identifier (ID) of  
13 each datagram; and

14 at least one network adapter coupled to the memory and processor to transmit the  
15 fragment over the fragment's associated network link.

1 28. (PREVIOUSLY PRESENTED) A system for distributing data over a plurality of  
2 network links within a computer network, comprising:

3 means for defining a unit of data as a datagram;

4 means for apportioning each datagram into at least one fragment;

5 means for associating each fragment to a network link of the plurality of network  
6 links according to a round robin policy based at least in part on an Internet protocol (IP)  
7 identifier (ID) of each datagram;

8 means for transmitting the fragment over the fragment's associated network link.

1 29. (PREVIOUSLY PRESENTED) A computer readable medium containing executable  
2 program instructions for execution on a processor, the executable program instructions  
3 comprising program instructions for:

4 defining a unit of data as a datagram;

5 apportioning each datagram into at least one fragment;

6 associating each fragment to a network link of the plurality of network links ac-  
7 cording to a round robin policy based at least in part on an Internet protocol (IP) identifier  
8 (ID) of each datagram;

9 transmitting the fragment over the fragment's associated network link.

1 30. (PREVIOUSLY PRESENTED) A method for distributing data over a plurality of  
2 network links within a computer network, comprising the steps of:

3 dividing a first datagram, having a first Internet protocol (IP) identifier (ID), into  
4 one or more fragments, each fragment of the first datagram associated with the first  
5 Internet protocol (IP) identifier (ID);

6 selecting a first network link of the plurality of network links for transmission of  
7 the one or more fragments of the first datagram;

8 transmitting all of the one or more fragments associated with the first IP ID over  
9 the first network link;

10           dividing a second datagram, having a second Internet protocol (IP) identifier (ID),  
11   into one or more fragments, each fragment of the second datagram associated with the  
12   second Internet protocol (IP) identifier (ID);

13           selecting a second network link of the plurality of network links for transmission  
14   of the one or more fragments of the second datagram; and

15           transmitting all of the one or more fragments associated with the second IP ID  
16   over the second network link.

1   31. (PREVIOUSLY PRESENTED) The method of claim 30 wherein the first network  
2   link and the second network link are selected according to a round robin policy based on  
3   the IP ID of each datagram.

1   32. (CURRENTLY AMENDED) A system for distributing data over a plurality of net-  
2   work links within a computer network comprising:

3           a processor;

4           a memory coupled to the processor and having locations accessible by the proces-  
5   sor;

6           an operating system resident in the memory and executed by the processor, the  
7   operating system configured to implement a load balancing technique that divides a first  
8   datagram, having a first Internet protocol (IP) identifier (ID), into one or more fragments,  
9   each fragment of the first datagram associated with the first Internet protocol (IP) identi-  
10   fier (ID), ~~set~~select a first network link of the plurality of network links for transmission  
11   of the one or more fragments of the first datagram, transmit all of the one or more frag-  
12   ments associated with the first IP ID over the first network link, divide a second data-  
13   gram, having a second Internet protocol (IP) identifier (ID), into one or more fragments,  
14   each fragment of the second datagram associated with a second Internet protocol (IP)

15 identifier (ID), select a second network link of the plurality of network links for transmis-  
16 sion of the one or more fragments of the second datagram, and transmit all of the one or  
17 more fragments associated with the second IP ID over the second network link.

1 33. (PREVIOUSLY PRESENTED) The system of claim 32 wherein the first network  
2 link and the second network link are selected according to a round robin policy based on  
3 the IP ID of each datagram.

1 34. (NEW) The method of claim 1 wherein the step of associating further comprises ap-  
2 portioning data substantially equally over the plurality of underlying links of the aggre-  
3 gate within the computer network.

1 35. (NEW) The system of claim 10 wherein the virtual interface process is configured to  
2 apportion data substantially equally over the plurality of underlying links of the aggregate  
3 within the computer network.

1 36. (NEW) The system of claim 11 wherein the means for associating is configured to  
2 apportion data substantially equally over the plurality of underlying links of the aggregate  
3 within the computer network.

1 37. (NEW) The system of claim 27 wherein the round robin policy is further based, at  
2 least in part, on a number of network links in the plurality of links.

1 38. (NEW) The system of claim 27 wherein the virtual interface process is configured to  
2 produce a result representing a remainder by dividing the IP ID by the number of network  
3 links.

1 39. (NEW) The method of claim 38 wherein the virtual interface process is further con-  
2 figured to calculate the IP ID of each datagram in a sequential manner and to rotate the  
3 fragments of each datagram among all the network links of the plurality of network links  
4 to thereby ensure that all fragments having the same IP ID are provided to the same net-  
5 work link.

1 40. (NEW) The system of claim 27 wherein the virtual interface process is configured to  
2 logically combine the IP ID with a predetermined mask to produce a quantity, right shift  
3 the quantity by a predetermined number of places to create a new quantity, and establish  
4 a threshold at which a group of data is forwarded to each network link of the plurality of  
5 network links.

1 41. (NEW) The system of claim 40 wherein the virtual interface process is further con-  
2 figured to produce a result representing a remainder upon dividing the new quantity by a  
3 number of network links in the plurality of network links.